

Jet User Meeting

16 June 2010

Overview

- State of Jet
- HFIP Upgrade
- wJet Upgrade
- Transition Issues
 - OS Upgrade
 - Real-Time Experiments
- Data Transfer Issues
- Filesystem quotas
- HSMS Capacity Issues
- Future NOAA RDHPCS Computing
- Feedback

HFIP Upgrade

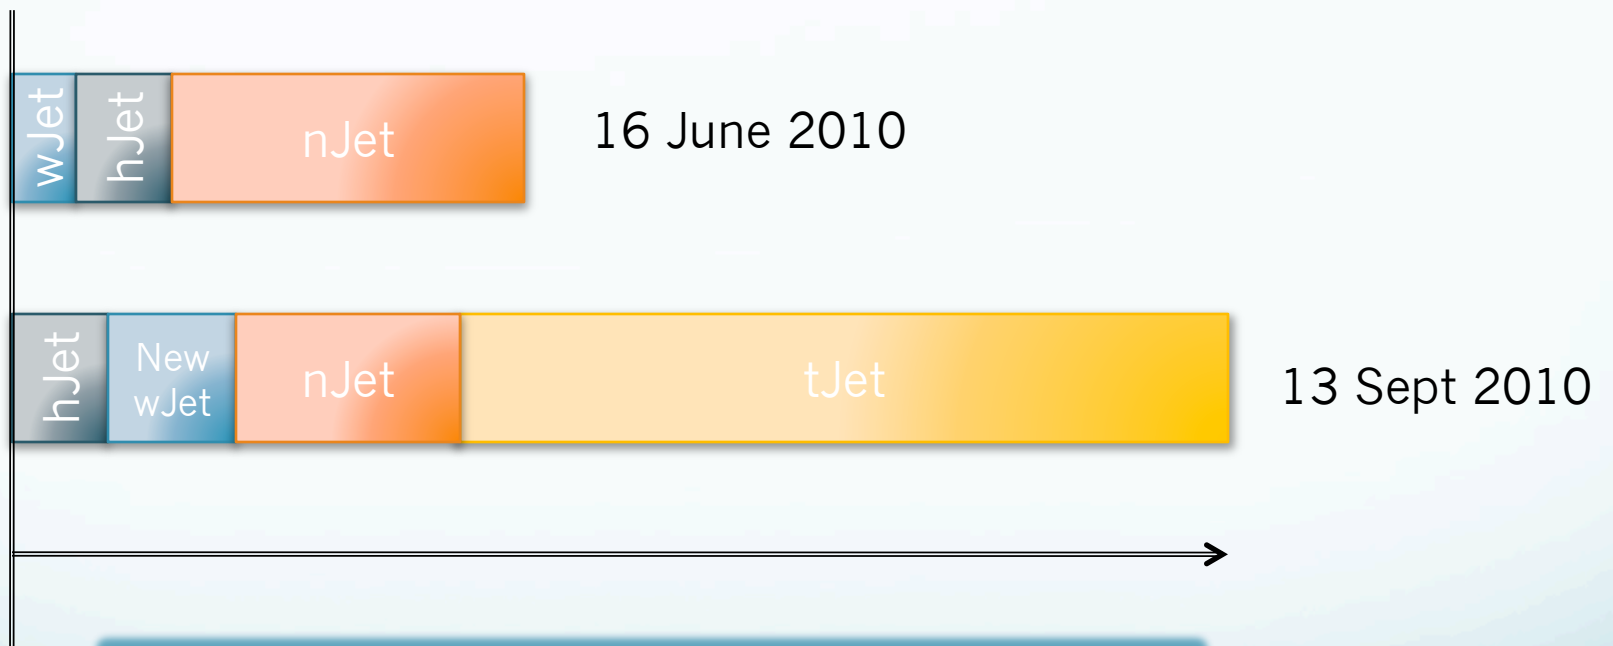
- Compute/storage upgrade for the Hurricane Forecast Improvement Project System
- Current system
 - ~2400 Intel Nehalem Cores
- Compute Upgrade
 - Repurpose approximately 1900 cores to other project
 - New 10,000 core Intel Westmere system to be available ~Aug 15th
- Storage Upgrade
 - Add ~100 TB to /lfs1

wJet Upgrade

- wJet is being decommissioned by Sept 13th
- wJet will be replaced by repurposed ~1376 HFIP cores
 - HFIP core is approximately 1.7x faster than wJet
- hJet will remain
- Add 76 TB of storage for wJet usage (on LFS1)

We will be providing more details and documentation on using the new wJet over the next month.

Summary of Systems at GSD



Size of box is relative to system performance
Bigger represents more performance

Total performance ~160 TF

System Transition Issues

- To support wJet decommission changes need to made to how the system operates
 - OS Upgrade
 - Real-time experiment migration

OS Upgrade

- The operating system on nJet will undergo minor changes and upgrades
 - Users should not have to recompile codes
- The new OS image will be installed on all systems
 - This will bring new compilers, new libraries and new tools
 - Simplify maintenance for the system team, which means we can bring tool updates faster
 - Allow more transparent use of both HFIP and Jet resources
 - All projects on wJet and hJet **will have** to be recompiled

Real-time experiment migration

- Current real-time experiments running on wJet (most all) will have to be moved
 - Will also have to be recompiled and validated on the new OS image
 - Any experiments running in on dedicated nodes much be migrated to reservations
 - Deadline of transition is September 13th, no exceptions

Schedule

Date	Event	Duration
16 June	Relocate /pan1	8 hours
17 June	Upgrade HSM	8 hours
25 June	Facilities Outage	1 hour
30 June	Replace /lfs1 controller	8 hours
14 July	Upgrade nJet OS New Service Nodes Begin hJet Upgrade	
26 July	Major downtime	2 days
8 Aug	/lfs1 upgrade	24 hours
27-29 Aug	Benchmarking and System Testing	3 days
1 Sept	**Decommission wJet Front Ends	
13 Sept	Decommission wJet	

** Jobs can still be submitted to wJet, but you cannot recompile applications for use on wJet.

Data Transfer Issues

- There have been ongoing data transfer issues affecting performance in and out of Jet
 - Affects long distance transfers, and in particular transfers to/from NCEP
 - Only fix is new hardware (which is on order)
 - Workaround is to use bbcp to the copy-hosts

BBCP

- BBCP is a parallel transfer that uses ssh for its authentication mechanism
 - Easy to get going
 - Measured performance:
 - > 50 MB/s to/from TACC
 - > 75 MB/s to/from NCAR
- Often installed on at other major HPC sites
 - At NCEP: /usrx/local/bbcp/bin/rs_aix4w2/bbcp
 - Up to 10 MB/s per bbcp instance
 - We have seen > 40 MB/s aggregate to/from NCEP
- For more details, email the help system

Filesystem Quotas

- Usage on lfs0 and lfs1 continues to increase
 - They are mostly full
 - Many accounts are grossly exceeding their quotas

Policy:

Accounts exceeding their quota will be warned by email. Users will have one week to reduce their disk usage. After one week, the account's ability to run jobs will be suspended until disk usage is reduced.

HSMS Capacity Issues

- We are running out of capacity on the HSM
 - Currently 1PB data is store on the archive
 - Should increase by 3x by Jan 2012
- We are working plans to resolve the issue
- Data older than 1 year may be shelved
 - It will be available, access will not be instantaneous
- Full solution implemented by Jan 2012
 - Will provide details when solution is determined

Future of NOAA HPC

- NOAA is moving towards centralized resources
 - Climate/Research System – Oak Ridge National Laboratory
 - First phase available late fall 2010 - Cray
 - Primarily for Climate Research, some resources will be available for other projects
 - Total system size by end of 2011 should be 1 PFlop
 - Weather/Production System – West Virginia
 - Most of the general HPC workload in Boulder will migrate to this system, early 2012
 - System should be at least 300 TF

Feedback

- Any thoughts, comments, concerns or questions?

Ask always, please ask any questions through the help system at:

hpcshelp.fsl@noaa.gov